

WHAT IS CLAIMED IS:

1 1. A method, comprising:
2 providing a source map indicating blocks of data striped across a first plurality of
3 storage units and a destination map indicating blocks of data striped across a second
4 plurality of storage units, wherein data is migrated from stripes indicated in the source
5 map to corresponding stripes indicated in the destination map; and
6 in response to determining that the source stripe and the destination stripe occupy
7 a same physical location on the storage units, writing the data from a source stripe to a
8 copy area and writing the data from the copy area to a corresponding destination stripe.

1 2. The method of claim 1, further comprising:
2 determining units of operation, wherein one unit of operation comprises one stripe
3 in the source map to migrate to one stripe in the second map, wherein the data is migrated
4 by processing the units of operation.

1 3. The method of claim 2, further comprising:
2 locking data in one source stripe in one unit of operation currently being
3 migrated; and
4 unlocking the locked data after completing the migration of the source data in the
5 unit of operation.

1 4. The method of claim 2, further comprising:
2 indicating a number of a current unit of operation being processed; and
3 indicating data is being copied through the copy area in response to determining
4 that the source stripe and destination stripe involved in the current unit of operation
5 occupy the same physical locations.

1 5. The method of claim 4, further comprising:
2 incrementing the current unit of operation in response to completing copying the
3 source stripe to the destination stripe for one unit of operation; and

4 indicating data is not being copied through the copy area in response to
5 completing copying the source stripe to the destination stripe for one unit of operation.

1 6. The method of claim 4, further comprising:
2 determining whether data is indicated as being copied through the copy area after
3 recovering from a failure;
4 writing the data from the copy area to the destination stripe in the indicated
5 current unit of operation in response to determining that the data is indicated as being
6 copied through the copy area; and
7 continuing processing the units of operation to complete the migration of the
8 source stripes to the destination stripes.

1 7. The method of claim 1, further comprising:
2 determining a depth of a source volume including the source stripes and a depth
3 of a destination volume including the destination the destination stripes;
4 writing the source stripes in descending order from one source stripe at a first
5 physical location of the source volume to the destination stripes in response to
6 determining that the destination volume depth exceeds the source volume depth; and
7 writing the source stripes in ascending order from one source stripe at a last
8 physical location of the source volume to the destination stripes in response to
9 determining that the destination volume depth does not exceed the source volume depth.

1 8. The method of claim 1, further comprising:
2 determining whether the migration of the data in the source map to locations in
3 the destination map is impermissible; and
4 aborting the migration in response to determining that the overlap is
5 impermissible.

1 9. The method of claim 8, wherein the migration is permissible comprises:
2 determining a depth of a source volume including the source stripes and a depth
3 of a destination volume including the destination stripes;

4 determining a source physical location on one storage unit of a first block in a
5 first stripe in the destination volume and a destination physical location on one storage
6 unit of a first block in a first stripe in the source volume; and
7 determining that the migration is impermissible in response to determining: (1)
8 that the destination volume depth is less than or equal to the source volume depth and the
9 destination physical location is greater than the source physical location or (2) that the
10 destination volume depth is greater than the source volume depth and the destination
11 physical location is less than the source physical location.

1 10. The method of claim 1, wherein a number of the first plurality of storage
2 units is different than a number of the second plurality of storage units.

1 11. The method of claim 1, further comprising:
2 detecting a failure of one of the first plurality of storage units;
3 rebuilding data from the failed storage units using parity data on the first plurality
4 of storage units that did not fail, wherein the migration is performed to transfer the rebuilt
5 data and the data in the first plurality of storage units that did not fail to the second
6 plurality of storage units including the storage units of the first plurality that survived.

1 12. The method of claim 1, wherein the storage units comprise magnetic hard
2 disk drives and wherein a Redundant Array of Independent Disk (RAID) algorithm is
3 used to stripe the data to the disks.

1 13. A storage controller in communication with a plurality of storage units,
2 comprising:
3 a source map indicating blocks of data striped across a first plurality of storage
4 units;
5 a destination map indicating blocks of data striped across a second plurality of
6 storage units; and
7 circuitry operable to:

8 (i) migrate data from stripes indicated in the source map to corresponding
9 stripes indicated in the destination map; and
10 (ii) in response to determining that the source stripe and the destination
11 stripe occupy a same physical location on the storage units, write the data from a
12 source stripe to a copy area and writing the data from the copy area to a
13 corresponding destination stripe.

1 14. The system of claim 13, wherein the circuitry is further operable to:
2 determine units of operation, wherein one unit of operation comprises one stripe
3 in the source map to migrate to one stripe in the second map, wherein the data is migrated
4 by processing the units of operation.

1 15. The system of claim 14, wherein the circuitry is further operable to:
2 lock data in one source stripe in one unit of operation currently being migrated;
3 and
4 unlock the locked data after completing the migration of the source data in the
5 unit of operation.

1 16. The system of claim 13, wherein the circuitry is further operable to:
2 indicate a number of a current unit of operation being processed; and
3 indicate data is being copied through the copy area in response to determining that
4 the source stripe and destination stripe involved in the current unit of operation occupy
5 the same physical locations.

1 17. The system of claim 16, wherein the circuitry is further operable to:
2 increment the current unit of operation in response to completing copying the
3 source stripe to the destination stripe for one unit of operation; and
4 indicate data is not being copied through the copy area in response to completing
5 copying the source stripe to the destination stripe for one unit of operation.

1 18. The system of claim 16, wherein the circuitry is further operable to:
2 determine whether data is indicated as being copied through the copy area after
3 recovering from a failure;
4 write the data from the copy area to the destination stripe in the indicated current
5 unit of operation in response to determining that the data is indicated as being copied
6 through the copy area; and
7 continue processing the units of operation to complete the migration of the source
8 stripes to the destination stripes.

1 19. The system of claim 13, wherein the circuitry is further operable to:
2 determine a depth of a source volume including the source stripes and a depth of a
3 destination volume including the destination the destination stripes;
4 write the source stripes in descending order from one source stripe at a first
5 physical location of the source volume to the destination stripes in response to
6 determining that the destination volume depth exceeds the source volume depth; and
7 write the source stripes in ascending order from one source stripe at a last physical
8 location of the source volume to the destination stripes in response to determining that the
9 destination volume depth does not exceed the source volume depth.

1 20. The system of claim 13, wherein the circuitry is further operable to:
2 determine whether the migration of the data in the source map to locations in the
3 destination map is impermissible; and
4 abort the migration in response to determining that the overlap is impermissible.

1 21. The system of claim 20, wherein the circuitry is further operable to:
2 determine a depth of a source volume including the source stripes and a depth of a
3 destination volume including the destination stripes;
4 determine a source physical location on one storage unit of a first block in a first
5 stripe in the destination volume and a destination physical location on one storage unit of
6 a first block in a first stripe in the source volume; and

7 determine that the migration is impermissible in response to determining: (1) that
8 the destination volume depth is less than or equal to the source volume depth and the
9 destination physical location is greater than the source physical location or (2) that the
10 destination volume depth is greater than the source volume depth and the destination
11 physical location is less than the source physical location.

1 22. The system of claim 13, wherein a number of the first plurality of storage
2 units is different than a number of the second plurality of storage units.

1 23. The system of claim 13, wherein the circuitry is further operable to:
2 detect a failure of one of the first plurality of storage units;
3 rebuild data from the failed storage units using parity data on the first plurality of
4 storage units that did not fail, wherein the migration is performed to transfer the rebuilt
5 data and the data in the first plurality of storage units that did not fail to the second
6 plurality of storage units including the storage units of the first plurality that survived.

1 24. A system, comprising:
2 a plurality of storage units;
3 a source map indicating blocks of data striped across a first plurality of storage
4 units;
5 a destination map indicating blocks of data striped across a second plurality of
6 storage units; and
7 circuitry operable to:
8 (i) migrate data from stripes indicated in the source map to corresponding
9 stripes indicated in the destination map; and
10 (ii) in response to determining that the source stripe and the destination
11 stripe occupy a same physical location on the storage units, write the data from a
12 source stripe to a copy area and writing the data from the copy area to a
13 corresponding destination stripe.

1 25. The system of claim 24, wherein the circuitry is further operable to:
2 indicate a number of a current unit of operation being processed; and
3 indicate data is being copied through the copy area in response to
4 determining that the source stripe and destination stripe involved in the current
5 unit of operation occupy the same physical locations.

1 26. The system of claim 24, wherein the storage units comprise magnetic hard
2 disk drives and wherein a Redundant Array of Independent Disk (RAID) algorithm is
3 used to stripe the data to the disks.

1 27. An article of manufacture capable of causing operations, the operations
2 operable to:
3 provide a source map indicating blocks of data striped across a first plurality of
4 storage units and a destination map indicating blocks of data striped across a second
5 plurality of storage units, wherein data is migrated from stripes indicated in the source
6 map to corresponding stripes indicated in the destination map; and
7 in response to determining that the source stripe and the destination stripe occupy
8 a same physical location on the storage units, write the data from a source stripe to a copy
9 area and writing the data from the copy area to a corresponding destination stripe.

1 28. The article of manufacture of claim 27, wherein the operations are further
2 operable to:
3 determine units of operation, wherein one unit of operation comprises one stripe
4 in the source map to migrate to one stripe in the second map, wherein the data is migrated
5 by processing the units of operation.

1 29. The article of manufacture of claim 28, wherein the operations are further
2 operable to:
3 lock data in one source stripe in one unit of operation currently being migrated;
4 and

5 unlock the locked data after completing the migration of the source data in the
6 unit of operation.

1 30. The article of manufacture of claim 28, wherein the operations are further
2 operable to:
3 indicate a number of a current unit of operation being processed; and
4 indicate data is being copied through the copy area in response to determining that
5 the source stripe and destination stripe involved in the current unit of operation occupy
6 the same physical locations.

1 31 The article of manufacture of claim 31, wherein the operations are further
2 operable to:
3 increment the current unit of operation in response to completing copying the
4 source stripe to the destination stripe for one unit of operation; and
5 indicate data is not being copied through the copy area in response to completing
6 copying the source stripe to the destination stripe for one unit of operation.

1 32. The article of manufacture of claim 30, wherein the operations are further
2 operable to:
3 determine whether data is indicated as being copied through the copy area after
4 recovering from a failure;
5 write the data from the copy area to the destination stripe in the indicated current
6 unit of operation in response to determining that the data is indicated as being copied
7 through the copy area; and
8 continue processing the units of operation to complete the migration of the source
9 stripes to the destination stripes.

1 33. The article of manufacture of claim 27, wherein the operations are further
2 operable to:
3 determine a depth of a source volume including the source stripes and a depth of a
4 destination volume including the destination the destination stripes;

5 write the source stripes in descending order from one source stripe at a first
6 physical location of the source volume to the destination stripes in response to
7 determining that the destination volume depth exceeds the source volume depth; and
8 write the source stripes in ascending order from one source stripe at a last physical
9 location of the source volume to the destination stripes in response to determining that the
10 destination volume depth does not exceed the source volume depth.

1 34. The article of manufacture of claim 27, wherein the operations are further
2 operable to:
3 determine whether the migration of the data in the source map to locations in the
4 destination map is impermissible; and
5 abort the migration in response to determining that the overlap is impermissible.

1 35. The article of manufacture of claim 34, wherein the operations are further
2 operable to:
3 determine a depth of a source volume including the source stripes and a depth of a
4 destination volume including the destination stripes;
5 determine a source physical location on one storage unit of a first block in a first
6 stripe in the destination volume and a destination physical location on one storage unit of
7 a first block in a first stripe in the source volume; and
8 determine that the migration is impermissible in response to determining: (1) that
9 the destination volume depth is less than or equal to the source volume depth and the
10 destination physical location is greater than the source physical location or (2) that the
11 destination volume depth is greater than the source volume depth and the destination
12 physical location is less than the source physical location.

1 36. The article of manufacture of claim 27, wherein a number of the first
2 plurality of storage units is different than a number of the second plurality of storage
3 units.

1 37. The article of manufacture of claim 27, wherein the operations are further
2 operable to:
3 detect a failure of one of the first plurality of storage units;
4 rebuild data from the failed storage units using parity data on the first plurality of
5 storage units that did not fail, wherein the migration is performed to transfer the rebuilt
6 data and the data in the first plurality of storage units that did not fail to the second
7 plurality of storage units including the storage units of the first plurality that survived.

1 38. The article of manufacture of claim 27, wherein the storage units comprise
2 magnetic hard disk drives and wherein a Redundant Array of Independent Disk (RAID)
3 algorithm is used to stripe the data to the disks.